Lesotho National Artificial Intelligence Policy

Green Paper (Early-Stage Policy Discussion)

Country: Lesotho

Date: October 13, 2025

Author: Norbert Butare

National Artificial Intelligence Green Paper (Early-Stage Policy Discussion Draft)

National Artificial Intelligence Green Paper (Early-Stage Policy Discussion Draft) **1**
- Executive Overview 2
- 1. Scope and Definitions 3
- 2. Policy Rationale & Context 4
- 3. Guiding Principles (Policy-Level) 5
- 4. Risk & Trust Framework (Conceptual) 6
- 5. Regulatory Posture: Options & Trade-offs (No Recommendations) 8
- 6. Institutional Landscape (Descriptive Only) 9
- 7. Data Governance at Policy Level 10
- 8. Compute, Infrastructure & Ecosystem Considerations (Conceptual) 11
- 9. Rights, Ethics, and Inclusion (Principle Statements) 12
- 10. Cross-Sector Considerations (Descriptive Notes, Not Use-Case Plans) 13
- 11. International Cooperation & Standards (Neutral Mapping) 14
- 12. Open Questions for Consultation 15
- Annex A: Glossary of Terms 16
- Annex B: Illustrative Policy References 17
  - Closing Note 17

Executive Overview

Purpose of this Green Paper

This discussion frames early policy choices for how Lesotho can approach artificial intelligence (AI). It is intended to support Cabinet-level reflection and stakeholder dialogue prior to any decisions being made. The focus is explanatory and principle-based. It outlines concepts, options, and trade-offs rather than providing instructions.

Why AI matters for Lesotho

AI is a general-purpose capability that may influence how the government delivers services across challenging terrain, how small businesses access markets, how farmers manage weather variability, and how clinics and schools augment scarce human expertise. AI may also introduce risks to rights, safety, and social cohesion—especially where data are limited, connectivity is uneven, or governance capacity is stretched. Economic shifts could bring new opportunities in digital services and entrepreneurship, but also raise concerns about job transitions, market power concentration, and reliance on foreign technologies.

What this document does and does not do

This Green Paper explains the scope and definitions, articulates high-level guiding principles, sketches a conceptual risk and trust framework, and neutrally maps regulatory posture options and institutional arrangements. It also outlines policy-level considerations for data governance, compute and infrastructure, rights and inclusion, and cross-sector opportunities and risks. It invites open questions for consultation.

It does not provide an implementation plan, detailed rules, costs, targets, timelines, ownership assignments, or step-by-step actions.

1. Scope and Definitions

Working definitions (Industry references; they do not create legal meaning):

- **AI system:** A system designed to process data and generate outputs such as predictions, classifications, recommendations, or content, using techniques that include machine learning, rules-based logic, statistical methods, or combinations thereof.
- **Model developer:** An entity or person that creates, trains, or substantially modifies an AI model (including pre-training and fine-tuning).
- **Deployer:** An organization or person that integrates an AI system into a use case and makes it available to end-users or uses it to make or support decisions.
- **Operator:** A person who configures or runs an AI system in a particular context (e.g., a civil servant using a decision-support tool).
- Affected individuals. Any person whose rights, opportunities, access to services, or safety could be impacted by an AI system's output or the decisions it informs.
- **High-risk context:** A setting where an AI system's failure or misuse could foreseeably cause significant harm to safety, fundamental rights, access to essential services, or democratic processes.
-

Boundaries of policy scope

For purposes of discussion, the policy scope could cover:

- **In scope:** Development, deployment, and use of AI systems by public bodies and private entities operating in Lesotho or delivering services to people in Lesotho; procurement and use of AI-enabled products; cross-border provision where outputs materially affect people in Lesotho.
- **Intersecting domains:** Data governance, cloud/compute access, connectivity, cybersecurity, and consumer protection—each of which influences AI outcomes.
- **Out of scope (for now):** Purely personal/household uses that do not affect others; defense and national security applications where separate regimes may apply; research prototypes not made available to users or decision-makers.

Emerging technology intersections.

AI policy may be affected by the maturity of identity systems, payment rails, e-government platforms, and spectrum and connectivity policies. Cloud and edge computing choices may shape access and resilience. Data policies (e.g., open data, localization, sharing) may influence innovation while raising privacy and security considerations.

2. Policy Rationale & Context

National context (conceptual)

Lesotho's economic base includes agriculture, textiles, water resources, public services, and remittances. Geography and settlement patterns create service-delivery challenges; connectivity is improving but uneven. The public sector faces capacity constraints; small firms and young entrepreneurs seek pathways to regional and global markets. In this context, AI could help extend scarce expertise (e.g., in health, education, justice administration), support climate adaptation in agriculture, improve citizen service channels, and enhance infrastructure maintenance. The issue is how to encourage beneficial experimentation while safeguarding rights and maintaining public trust.

Regional and international alignment landscape (descriptive)

Lesotho operates within the African Union (AU) and Southern African Development Community (SADC) ecosystems and engages with multilateral bodies such as UNESCO, OECD, and the International Telecommunication Union (ITU). These fora have issued high-level principles, ethics guidance, and standardization workstreams on AI. Alignment and mutual recognition could reduce fragmentation, support cross-border data flows, and facilitate trade—while still reflecting local context and constitutional values.

3. Guiding Principles (Policy-Level)

The following principles are framed as policy intent. They do not prescribe controls. They indicate directions that subsequent deliberations could translate into appropriate instruments.

1. **Lawfulness.**
   AI uses should be consistent with the Constitution and applicable laws (e.g., privacy, consumer, competition).
   Policy implication: Future measures could clarify how existing laws apply to AI and identify any gaps that may exist.

2. **Human rights and dignity.**
   AI should respect rights to equality, privacy, freedom of expression, and due process.
3. ***Policy implication:*** High-risk deployments should consider incorporating human oversight and safeguards for rights-impacting decisions.

4. **Safety and security.**
   AI systems should aim to avoid foreseeable harm to people and critical infrastructure, including cybersecurity risks.
   Policy implication: Safety testing and secure-by-design practices may be encouraged, particularly in higher-risk contexts.

5. **Privacy and data protection.**
   Personal data use in AI should reflect lawful bases, purpose limitation, data minimization, and appropriate safeguards.
   Policy implication: Data governance instruments could clarify expectations on collection, re-use, and sensitive data.

6. **Accountability.**
   Responsibility for AI outcomes should be traceable across the lifecycle (developer, deployer, operator).
   Policy implication: Clear documentation and well-defined lines of accountability can support effective oversight and redress.

7. **Transparency and explainability (proportionate).**
   People may benefit from knowing when AI is involved and receiving understandable explanations where decisions significantly affect them.
   Policy implication: Proportionate disclosure and explanation could be considered for rights-affecting uses.

8. **Fairness and non-discrimination.**
   AI should seek to avoid unjust bias, especially in the allocation of public benefits and access to services.
   Policy implication: Context-appropriate testing for disparate impacts may be explored.

9. **Accessibility and inclusion.**
   AI should consider language, literacy, disability, and connectivity barriers, and reflect cultural and linguistic context (including Sesotho).
   Policy implication: Accessibility features and local language support may be encouraged.

10. **Environmental stewardship.**
    AI development and use should consider energy and water impacts, and support climate resilience.
    Policy implication: Efficiency, shared compute, and low-resource model options could be examined.
11. **Proportionality.**
    Measures should be commensurate with risk and capacity, avoiding undue burdens on innovation or on small organizations.
    Policy implication: Scaled expectations could support SMEs and early experimentation.

12. **Interoperability and openness.**
    Open standards and interoperable approaches may promote competition, portability, and resilience.
    Policy implication: Standards adoption could be a focus for alignment rather than bespoke technical rules.

13. **Public benefit and trust.**
    The goal is socially beneficial AI that sustains confidence in public institutions.
    Policy implication: Mechanisms for public engagement and feedback may be considered.

4. Risk & Trust Framework (Conceptual)

Risk taxonomy (illustrative)

- Minimal risk: Tools with negligible impact on people's rights or safety, such as grammar suggestions or non-critical data summaries for internal use.
  . *Illustrative example:* An office productivity assistant recommending document formatting.

- Limited risk: Systems interacting with users but not making consequential decisions, where simple disclosures can mitigate concerns.
  Illustrative example: A chatbot providing general information about citizen services with a clear hand-off to human agents.

- High-risk: Uses that could significantly affect access to essential services, allocate public benefits, influence credit or employment, mediate health or safety decisions, or impact the justice system.
  Illustrative examples: Triage support in clinics; targeting for social assistance; risk scoring in school placement or scholarships.

- Prohibited (for consideration). Uses that would be fundamentally incompatible with rights or safety (e.g., manipulative or exploitative practices targeting vulnerable groups; indiscriminate biometric surveillance in public spaces).
  Note: The boundary of any prohibited category would require careful consultation and legal analysis.

Trust measures in principle:

To support confidence without prescribing methods, the following measures may be considered conceptually:

- **Documentation**: Model and system summaries (e.g., data provenance, intended uses, limitations).
- **Testing and evaluation:** Pre-deployment and ongoing assessment of performance, robustness, fairness, and security.
- **Oversight and human-in-the-loop:** Appropriate human judgment for consequential decisions.
- Incident handling: Processes for recording and learning from adverse events or near misses.
- **Transparency to users:** Proportionate disclosures that AI is involved; accessible explanations for significant effects.
- Redress: Channels for complaints, appeals, and correction of errors.
- Independent assurance: Where helpful, use of audits, certifications, or conformity assessments referencing recognized standards.

5. Regulatory Posture: Options & Trade-offs (No Recommendations)

Option A: Sector-led guidance with central coordination.

- Description: Sector ministries or regulators develop AI guidance aligned to their mandates, with a central unit providing coherence and shared tools.
- Pros: Leverages domain expertise; flexible and iterative; can align with existing sector laws.

- Cons: Risk of fragmentation; uneven protections; coordination overhead; potential regulatory gaps for cross-cutting risks.

Option B: Central AI statute with risk-based provisions.

- Description: A cross-cutting AI law articulating definitions, risk categories, and baseline duties, complemented by sector rules.
- Pros: Coherent baseline; more explicit rights protections; easier international signaling and mutual recognition.
- Cons: Risk of rigidity; drafting complexity; capacity demands for enforcement; potential duplication with existing laws.

Option C: Co-regulation and standards-first model.

- Description: Government articulates principles and outcomes; recognized standards and codes of practice provide technical detail; compliance may be demonstrated through attestations or certifications.
- Pros: Keeps pace with technology; taps international standards; lowers compliance cost for SMEs.
- Cons: Dependence on standards development; questions about accountability; standards may not fully reflect the local context.

Option D: Sandbox-and-learning posture at an early stage.

- Description: Structured environments where organizations test AI uses under agreed safeguards, to generate evidence for future policy.
- Pros: Builds capacity; surfaces real-world risks; informs proportionate rules.
- Cons: Limited scale; potential perceptions of favoritism; requires careful guardrails to protect rights.

6. Institutional Landscape (Descriptive Only)

Possible roles (illustrative):

- Policy ministry: Articulates national AI policy intent and coordinates across government.
- Independent regulator (generic): Oversees compliance with cross-cutting AI or digital market rules if such a mandate emerges.
- Central digital/AI unit: Provides guidance, templates, and shared resources; convenes learning across sectors.
- Data protection authority: Oversees personal data issues relevant to AI, including individual rights and lawful processing.
- Standards body and accreditation actors: Promote adoption of recognized standards, testing, and conformity assessment.

- Ethics advisory forum: Creates a channel for public, academic, and civil society input on sensitive or novel AI uses.

Coordination models and accountability concepts (conceptual):

- **RACI-style logic:** For each significant AI use, there could be clarity about who is *Responsible* for development/deployment, who is *Accountable* for outcomes, who is *consulted* (e.g., data protection, security), and who is *informed* (e.g., the public).
- **Escalation pathways:** Complex or high-risk uses could be escalated to a cross-functional review.
- **Independence vs. integration:** An independent oversight function may enhance trust; integrated sector oversight may ensure context sensitivity. The balance is a policy choice.

7. Data Governance at Policy Level

Data stewardship:

Stewardship involves clear roles for how data are collected, curated, shared, and retired. Emphasis could be placed on data quality, context, and representativeness to avoid errors and bias in AI outputs.

Quality and provenance:

Provenance (where the data came from and how they were processed) supports accountability. Mechanisms for recording lineage and maintaining metadata could be encouraged, especially in high-impact datasets.

Access and sharing:

Policy choices include when to enable data access for innovation (e.g., research, SMEs) and when to restrict access to protect rights and commercial confidentiality. Trusted data-sharing arrangements may balance these interests.

Localization and interoperability:

Some stakeholders may prefer local storage for sensitivity or resilience; others may seek

regional cloud options for cost and performance. Interoperability and cross-border flow arrangements can reduce fragmentation while respecting legal protections.

Open data vs. protected data:
Open data can spur innovation and transparency, but not all data is suitable for open release. Sensitive data (e.g., health, justice, financial) may require strong safeguards or anonymization before sharing.

Privacy and security principles:
High-level principles include lawfulness, purpose limitation, data minimization, security by design, and user rights (access, correction, objection). Security considerations span encryption, access controls, logging, and incident handling—without prescribing tools.

Trade-offs:
Innovation benefits from access and re-use; protection requires limits and safeguards. The appropriate balance may vary by context and risk.

8. Compute, Infrastructure & Ecosystem Considerations (Conceptual)

National vs. regional compute access:
Advanced AI often depends on specialized hardware. Choices include limited national capacity, regional pooling, or commercial cloud access. Regional models offer scale and cost advantages, while national capacity can support sovereignty and resilience. A mixed approach is common internationally.

Cloud and edge considerations:
Cloud services can provide elasticity and access to state-of-the-art tools. Edge computing

(on-device or near the user) may reduce latency, costs, and data movement, benefiting remote areas with intermittent connectivity.

Connectivity dependencies:

Backbone and last-mile connectivity influence AI feasibility. Policies on spectrum, backhaul, and community networks interact with AI benefits, especially for rural service delivery.

Standards and assurance thinking:

Referencing recognized technical and management standards may reduce risk and improve interoperability. Conformity assessments and certifications—where appropriate—could provide signals of trust for high-impact systems.

Innovation ecosystem (as themes, not programs):

- Research and knowledge creation: Links between universities, public sector challenges, and industry may cultivate relevant expertise.
- Skills and workforce: Digital literacy, data skills, and responsible AI awareness across professions (health, education, agriculture, justice) may be foundational.
- Entrepreneurship and SME: Access to data, toolkits, and computing can lower entry barriers.
- Demand signals: Public procurement policies—designed around principles, not prescriptive specs—can encourage trustworthy solutions.

9. Rights, Ethics, and Inclusion (Principle Statements)

Human oversight:

Where AI influences consequential decisions, humans could retain meaningful oversight with the ability to question, override, or appeal results.

Non-discrimination:

AI should avoid unjust impacts on groups based on gender, disability, language, location, age, or socioeconomic status. Testing for disparate outcomes, alongside qualitative evaluation, may be informative.

Accessibility and language:

Systems could consider local language (Sesotho) and dialects, low-literacy interfaces, offline functionality, and accessibility for persons with disabilities.

Cultural and linguistic context:

Training data and content policies may reflect local cultural norms and context, avoiding stereotyping or cultural erasure.

Redress and grievance (in principle):

People affected by AI-enabled decisions could have channels to seek explanations, corrections, or appeals. Oversight bodies may consider aggregated insights from complaints to inform and improve their practices.

10. Cross-Sector Considerations (Descriptive Notes, Not Use-Case Plans)

Public services and administration:

- Opportunities: Case triage, digital assistants for citizen queries, and analytics for infrastructure maintenance.
- Risks: Over-reliance on automated outputs, exclusion of those without digital access, and opacity in decision pathways.

- Policy questions: What constitutes a “significant effect” on a citizen? What level of disclosure is appropriate? How should records be kept for audit?

Health

- Opportunities: Clinical decision support, telemedicine triage, supply chain optimization for medicines.
- Risks: Safety and efficacy concerns, biased datasets, privacy risks, and accountability in adverse events.
- Policy questions: What thresholds define “high-risk” in health? How should human oversight be structured? What evidence is sufficient for deployment in resource-constrained settings?

Education

- Opportunities: Personalized learning aids, teacher workload reduction, and language support.
- Risks: Academic integrity, inequities in access, inappropriate content exposure, and data protection for minors.
- Policy questions: What disclosures are needed when AI supports assessment? How to balance innovation with safeguarding learners’ rights?

Agriculture

- Opportunities: Weather-informed advisory, pest detection, market price insights.
- Risks: Overfitting to limited local data, misaligned incentives from proprietary models, and exclusion of smallholders without connectivity.
- Policy questions: What data stewardship arrangements support trusted advisory services? How to ensure the inclusion of low-connectivity farmers?

Justice and public safety

- Opportunities: Case management analytics, legal information retrieval, and language translation for court users.
- Risks: Bias in risk scoring, due process concerns, surveillance overreach, and cybersecurity.
- Policy questions: Which uses are unsuitable (potential prohibition)? What transparency and contestability are appropriate in rights-affecting decisions?

11. International Cooperation & Standards (Neutral Mapping)

Why alignment matters:

Compatibility with regional and global norms can reduce trade friction, support investment, and enable cross-border services. It may also ease participation in shared computing and research initiatives.

Relevant fora and instruments (descriptive, non-exhaustive):

- African Union and SADC: Emerging continental AI discussions; data protection and cybersecurity model instruments.
- UNESCO: Ethical guidance on AI.
- OECD: Principles on AI and frameworks for risk and system classification.
- ITU: Telecommunications and standardization activities with AI dimensions.
- ISO/IEC JTC 1/SC 42: International technical standards on AI concepts, terminology, management systems, and risk.
- Other communities Include Standards development organizations, academic networks, and public-interest coalitions focused on AI safety, fairness, and transparency.

Mutual recognition and adequacy:

Mechanisms for recognizing equivalent safeguards across borders can facilitate data flows and market access while maintaining protections. Engagement in standards processes may help ensure local realities are reflected.

12. Open Questions for Consultation

1. How should Lesotho define the policy boundary for “AI systems” in a way that is technology-neutral yet practically useful?
2. Which types of AI used in public services, if any, should be considered presumptively high-risk?
3. What transparency to users would be meaningful and proportionate across different risk levels?
4. Where are current laws (privacy, consumer, competition, procurement, cybersecurity) sufficient, and where are targeted clarifications needed?
5. What institutional configuration would best balance independence, expertise, and coordination for AI oversight?
6. What documentation and assurance practices are feasible for SMEs and start-ups without creating undue burden?
7. How should grievance and redress mechanisms work for AI-affected individuals, especially in rural or low-connectivity settings?
8. What approaches to data stewardship and provenance would improve data quality for AI while respecting privacy and commercial confidentiality?
9. Where should Lesotho position itself on data localization versus regional or cloud-based storage, and under what safeguards?
10. How can open data be expanded responsibly, and what categories of data should remain protected or shared only under controlled access?

Annex A: Glossary of Terms

- AI system: A system that processes data to produce outputs such as predictions, classifications, recommendations, or content.
- AI model: A parameterized computational representation (e.g., neural network, decision tree) trained on data to perform tasks.
- Foundation model: A broad-scope model trained on large datasets that can be adapted (fine-tuned) for varied tasks.
- Fine-tuning: Additional training to adapt a model to a specific domain or task.
- Developer: An entity or person that creates or modifies an AI model or system.
- Deployer: An organization or person that integrates and uses an AI system to deliver a service or support a decision.
- Operator: A person who runs or configures an AI system in practice.
- Affected individual: A person whose rights, opportunities, or safety could be impacted by an AI system.
- Bias (statistical): Systematic error that skews outcomes, potentially leading to unfair impacts.
- Explainability: The extent to which the internal mechanics or the rationale for an output can be understood.
- Human-in-the-loop (HITL): A design approach where humans review or can override AI outputs.
- Data controller/data processor: Roles related to personal data governance (definitions may align with privacy law).
- Provenance: Records of the origin, context, and transformations of data or content.
- Synthetic content: Content generated by AI that may resemble real data or media.
- Watermarking / content labeling: Techniques to signal that content is AI-generated.
- Compute: Processing power required for AI training or inference (CPUs, GPUs, specialized accelerators).
- Edge computing: Processing data near the source or on devices, reducing the need for backhaul.
- Model card/system card: Summaries describing an AI model or system's intended use, limitations, and evaluation.
- Impact assessment: Structured consideration of potential effects on rights, safety, and society.
- Red-teaming: Adversarial testing to uncover risks, vulnerabilities, or misuse pathways.
- Audit (AI): Independent evaluation of an AI system's conformity to stated criteria.
- Assurance: Evidence that a system meets certain expectations (through tests, documentation, or certifications).
- Interoperability: The ability of different systems to exchange and use information.
- Proportionality: Aligning measures with the level of risk and potential impact.

Annex B: Illustrative Policy References

- *African Union. African Union Convention on Cyber Security and Personal Data Protection (Malabo Convention).African Union, 27 June 2014. African Union.*

- *IEEE Standards Association. IEEE Std 7000-2021: IEEE Standard Model Process for Addressing Ethical Concerns During System Design. IEEE, 15 Sept. 2021.*

- *International Organization for Standardization/International Electrotechnical Commission. ISO/IEC 22989:2022 — Information Technology — Artificial Intelligence — Concepts and Terminology. ISO/IEC, 2022.*

- *International Telecommunication Union. Focus Group on AI for Health: Benchmarking and Evaluation Reports. ITU, 2023.*

- *National Institute of Standards and Technology (NIST). AI Risk Management Framework 1.0. NIST, Jan. 2023.*

- *Organization for Economic Co-operation and Development. OECD Principles on Artificial Intelligence. OECD, 2019.*

- *Organization for Economic Co-operation and Development. Framework for the Classification of AI Systems. OECD, 2022.*

- *Southern African Development Community. SADC Model Law on Computer Crime and Cybercrime. SADC, 2013.*

- *Southern African Development Community. SADC Model Law on Data Protection. SADC, 2013.*
- *UNESCO. Recommendation on the Ethics of Artificial Intelligence. UNESCO, 2021.*

Closing Note

This Green Paper is exploratory. It seeks to clarify choices rather than to decide them. Stakeholders are invited to respond to the open questions, share evidence and perspectives, and propose refinements to the principles and conceptual frameworks outlined here.